

WHAT IS CLAIMED IS:

1. In a data replication system having a primary computer system and a backup computer system, a method of lock-step replication of database updates that occurred in
5 the primary computer system to the backup computer system, the method comprising:
 within a first application executing on the primary computer system:
 performing and completing a first transaction on the primary computer system, the first transaction updating a first file in the primary computer system;
 in the primary computer system, upon completing the first transaction,
10 initiating a lockstep transaction that updates a second file in the primary computer system;
 and
 waiting to receive a predefined message prior to performing any further operations;
 sending audit records from the primary computer system to the backup computer
15 system, the sent audit records including audit records representing the updates to the first file by the first transaction and the updates to the second file by the lockstep transaction;
 receiving from the backup computer system confirmation that the audit records representing the updates to the first file by the first transaction and the updates to the second file by the lockstep transaction have been durably stored by the backup computer
20 system, and upon receiving said confirmation, sending the predefined message to the first application.
2. The method of claim 1 wherein the lockstep transaction is initiated by a procedure call made immediately upon completion of the first transaction.
3. The method of claim 1 wherein the first application performs an operation
25 dependent upon completion of the first transaction only after receiving the first predefined message.

- 30 4. The method of claim 1 further comprising:
 upon occurrence of a pre-determined event that terminates the lockstep transaction, initiating a second lockstep transaction that updates the second file in the primary computer system;
 after the second lockstep transaction is initiated, sending audit records from the
35 primary computer system to the backup computer system, the sent audit records including

audit records representing the updates to the second file by the another lockstep transaction;

after the second lockstep transaction is initiated, ignoring said confirmation that the audit records representing the updates to the first file by the first transaction and the

updates to the second file by the lockstep transaction have been durably stored by the backup computer system;

after the second lockstep transaction is initiated, receiving a second confirmation that the audit records representing the updates to the second file by the second lockstep transaction have been durably stored by the backup computer system, and upon receiving

said second confirmation, sending the predefined message to the first application.

5. In a data replication system having a primary computer system and a backup computer system, a method of lock-step replication of database updates that occurred in the primary computer system to the backup computer system, the method comprising:

initiating a lockstep transaction;

generating a lockstep audit record corresponding to the lockstep transaction, the lockstep audit record having a first transaction identifier;

storing the lockstep audit record in an audit trail;

reading audit records stored in the audit trail in a sequence in which the audit

records are stored;

transmitting the audit records to the backup computer system, wherein the backup computer system includes mechanism for safely storing the lockstep audit record and audit records preceding the lockstep audit record immediately upon receiving the lockstep audit record, the backup computer system further including mechanisms for transmitting a safe audit trail position of the lockstep audit record to the primary computer system after the lockstep audit record is safely stored;

receiving the safe audit trail position from the backup computer system;

checking whether the safe audit trail position corresponds to a lockstep transaction that is currently active; and

based on results of the checking step, indicating completion of the lockstep replication procedure.

6. The method of claim 5, further comprising:

storing the first transaction identifier at a first location of a pre-defined data

structure; and

during the reading step and upon encountering the first lockstep audit record,
extracting an audit trail position and a transaction identifier from the first lockstep audit
record;

storing the extracted audit trail position at a second location of the pre-defined data
5 structure; and

storing the extracted transaction identifier at a third location of the pre-defined data
structure.

7. The method of claim 6, wherein the checking step comprises:

10 comparing the safe audit trail position to the audit trail position stored at the
second location; and

comparing the transaction identifier stored at the first location and the transaction
identifier stored at the third location.

15 8. The method of claim 7, further comprising:

upon occurrence of an event that disrupts the lockstep replication procedure before
completion, performing another lockstep transaction, the another lockstep transaction
having a new transaction identifier; and

20 storing the new transaction identifier in the first location of pre-defined data
structure such that the checking step results in a mismatch between the transaction
identifier stored at the first location and the transaction identifier stored at the third
location.

9. The method of claim 5, further comprising:

25 pausing execution of an application program upon initiation of the lockstep
replication procedure; and

resuming execution of the application program upon completion of the lockstep
replication procedure.

30 10. The method of claim 5, wherein the transmitting step comprises transmitting at
least a subset of the audit records to the backup computer system in a message buffer, and
wherein the backup computer system is configured to return an audit trail position of a last
saved audit record as the safe audit trail position without ensuring the audit records of the
message buffer are durably stored unless the lockstep audit record is included in the
35 message buffer.

11. In a data replication system having a primary computer system and a backup computer system, a method of lock-step replication of database updates that occurred in the primary computer system to the backup computer system, the method comprising:

initiating a first lockstep replication procedure and performing a first update on a pre-determined file in the primary system, the first update being identified by a first unique transaction identifier;

storing the first unique transaction identifier in a pre-defined data structure in the primary system as a lockstep gateway transaction identifier (LockStep_Gateway_TID);

generating audit records that indicate database updates pertaining to database transactions performed on the primary system, the audit records further including a first lockstep audit record that is associated with the first update on the pre-determined file and that includes the first unique transaction identifier;

storing the audit records in an audit trail in the primary system;

extracting audit records from the audit trail for transmission to the backup computer system;

storing an audit trail position of the first update in the pre-defined data structure upon encountering the first lockstep audit record during the extracting step;

storing the first unique transaction identifier in the pre-defined data structure as a lockstep audit transaction identifier (LockStep_Audit_TID) upon encountering the first lockstep audit record during the extracting step;

transmitting the stream of audit records and a lock-step indicator to the backup computer system, wherein the lock-step indicator indicates a lockstep replication procedure has initiated, wherein the backup computer system is configured to ensure the stream of audit records are durably stored upon receiving the lock-step indicator, and wherein the backup computer system is configured to transmit to the primary computer system a safe position indicating the audit trail position of durably stored audit records upon receiving the lock-step indicator;

comparing the safe position returned by the backup computer system to the audit trail position stored in the pre-defined data structure; and

indicating completion of the lockstep replication procedure when the safe position is equal to or higher than the audit trail position stored in the pre-defined data structure, and when the lockstep gateway transaction identifier (LockStep_Gateway_TID) matches the lockstep audit transaction identifier (LockStep_Audit_TID).

12. The method of claim 11, further comprising:

pausing execution of an application program upon initiation of the lockstep replication procedure; and

resuming execution of the application program upon completion of the lockstep replication procedure.

5

13. The method of claim 11, further comprising:

upon occurrence of an event that disrupts the first lockstep replication procedure before completion, performing a second update on the pre-determined file in the primary system, the second update being identified by a second unique transaction identifier;

10

storing the second unique transaction identifier in the pre-defined data structure as the lockstep gateway transaction identifier (LockStep_Gateway_TID) in place of the first unique transaction identifier.

14. The method of claim 11, wherein the transmitting step comprises transmitting the stream of audit records to the backup computer system one buffer at a time, and wherein the backup computer system is configured to return an audit trail position of a last saved audit record as the safe position without ensuring the audit records of the buffer are durably stored unless the lockstep indicator is included in the buffer.

15

20

15. In a data replication system having a primary computer system and a backup computer system, a method of lock-step replication of database updates that occurred in the primary computer system to the backup computer system, the process comprising:

starting a lockstep replication procedure;

performing a first update on a pre-determined file in the primary system, the first

25

update being identified by a first unique transaction identifier;

storing the first unique transaction identifier in a pre-defined data structure in the primary system as a lockstep gateway transaction identifier (LockStep_Gateway_TID);

generating audit records that indicate database updates pertaining to database transactions performed on the primary system, the audit records further including a first

30

lockstep audit record that is associated with the first update on the pre-determined file and that includes the first unique transaction identifier;

storing the audit records in an audit trail in the primary system;

upon an occurrence of an event that disrupts operations of the primary computer system, performing the steps of:

performing a second update on the pre-determined file in the primary system, the second update being identified by a second unique transaction identifier,

replacing the first unique transaction identifier with the second unique transaction identifier in the pre-defined data structure,

- 5 generating a second lockstep audit record that is associated with the second update on the pre-determined file and that includes the second unique transaction identifier,

storing the second lockstep audit record in the audit trail;

extracting audit records from the audit trail for transmission to the backup

- 10 computer system;

concurrently with the extracting step, storing audit trail position of the first lockstep audit record in the pre-defined data structure upon encountering the first lockstep audit record and replacing the stored audit trail position with the audit trail position of the second lock step audit record upon encountering the second lockstep audit record;

- 15 concurrently with the extracting step, storing the first unique transaction identifier in the pre-defined data structure as a lockstep audit transaction identifier (LockStep_Audit_TID) upon encountering the first lockstep audit record and replacing the stored lockstep audit transaction identifier with the second unique transaction identifier upon encountering the second lockstep audit record;

- 20 transmitting the stream of audit records and a lock-step indicator to the backup computer system, wherein the lock-step indicator indicates a lockstep replication procedure has initiated, wherein the backup computer system is configured to ensure the stream of audit records are durably stored upon receiving the lock-step indicator, and wherein the backup computer system is configured to transmit to the primary computer system a safe position indicating the audit trail position of durably stored audit records upon receiving the lock-step indicator;

comparing the safe position returned by the backup computer system to the audit trail position stored in the pre-defined data structure; and

- 30 indicating completion of the lockstep replication procedure when the safe position is equal to or higher than the audit trail position stored in the pre-defined data structure, and when the lockstep gateway transaction identifier (LockStep_Gateway_TID) matches the lockstep audit transaction identifier (LockStep_Audit_TID).

16. The method of claim 15, further comprising:

pausing execution of an application program upon starting the lockstep replication procedure; and

resuming execution of the application program upon completion of the lockstep replication procedure.

17. The method of claim 15, wherein the transmitting step comprises transmitting the stream of audit records to the backup computer system one buffer at a time, and wherein the backup computer system is configured to return an audit trail position of a last saved audit record as the safe position without ensuring the audit records of the buffer are durably stored unless the lockstep indicator is included in the buffer.

18. A database replication system having a primary computer system and a backup computer system, the primary computer system configured to couple to a database, the primary computer system having an application program that performs database transactions on the database, the database replication system comprising:

a gateway configured to initiate a lockstep replication procedure and perform a first update on a pre-determined file in the database upon receiving a lockstep request from the application program, wherein the first update is identified by a first unique transaction identifier;

a TMF module configured to generate audit records that indicate database updates pertaining to database transactions performed on the primary system, wherein the audit records further include a first lockstep audit record that is associated with the first update on the pre-determined file and that includes the first unique transaction identifier, the TMF module further configured to store the audit records in an audit trail in the primary system;

an extractor configured to extract audit records from the audit trail for transmission to the backup computer system;

the extractor configured to store the first unique transaction identifier received from the gateway process in a pre-defined data structure in the primary system as a lockstep gateway transaction identifier (LockStep_Gateway_TID);

the extractor configured to store an audit trail position of the first update in the pre-defined data structure upon encountering the first lockstep audit record in the audit trail;

the extractor configured to store the first unique transaction identifier in the pre-defined data structure as a lockstep audit transaction identifier (LockStep_Audit_TID)

upon encountering the first lockstep audit record in the audit trail;

the extractor configured to transmit the stream of audit records and a lock-step indicator to the backup computer system, wherein the lock-step indicator indicates a lockstep replication procedure has initiated, wherein the backup computer system is configured to ensure the stream of audit records are durably stored upon receiving the lock-step indicator, and wherein the backup computer system is configured to transmit to the extractor a safe position indicating the audit trail position of durably stored audit records upon receiving the lock-step indicator;

the extractor configured to compare the safe position returned by the backup computer system to the audit trail position stored in the pre-defined data structure;

the extractor configured to communicate to the gateway a status of the lockstep replication procedure when the safe position is equal to or higher than the audit trail position stored in the pre-defined data structure, and when the lockstep gateway transaction identifier (LockStep_Gateway_TID) matches the lockstep audit transaction identifier (LockStep_Audit_TID); and

the gateway configured to generate a response to the lockstep request according to the status of the lockstep replication procedure.

19. The data replication system of claim 18, wherein execution of the application program pauses upon initialization of the lockstep replication procedure and wherein execution of the application program is configured to resume upon completion of the lockstep replication procedure.

20. The data replication system of claim 18, wherein the gateway is configured to perform a second update on the pre-determined file in the primary system upon occurrence of an event that disrupts the first lockstep replication procedure before completion, the second update being identified by a second unique transaction identifier.

21. The data replication system of claim 20, wherein the gateway is configured to replace the second unique transaction identifier in the pre-defined data structure as the lockstep gateway transaction identifier (LockStep_Gateway_TID) in place of the first unique transaction identifier.

22. The data replication system of claim 18, wherein the extractor transmits the stream of audit records to the backup computer system one buffer at a time, and wherein the backup computer system is configured to return an audit trail position of a last saved audit

record as the safe position without ensuring the audit records of the buffer are durably stored unless the lockstep indicator is included in the buffer.

10056145" 012202